Open Science Grid

OSG Production
Foundations for 2M+ Hours/Day

April 9, 2014

Rob Quick

With Help from
Shawn McKee and Chander Seghal

Once Upon a Time…

Phoenix, November 2003

Super Computing

Agenda

- OSG Networking
- Capturing Opportunistic Cycles
- OSG Operations
- OSG as a Community

OSG Networking Area

- OSG Networking was added at the beginning of OSG's second 5-year period in 2012
- The "Mission" is to have OSG become the network service data **source** for its constituents
  - Information about network performance, bottlenecks and problems should be easily available.
  - Should support our VOs, users and site-admins to find network problems and bottlenecks.
  - Provide network metrics to higher level services so they can make informed decisions about their use of the network *(Which sources, destinations for jobs or data are most effective?)*
- **Goal**: OSG hosts network information for its constituents, aiding in finding/fixing problems and enabling applications and users to better take advantage of their networks

Year 1&2 Goals and Key Initiatives in Network Area

- Year 1 of OSG Networking was primarily focused on getting network monitoring in place
  - **Deploying perfSONAR-PS:** Instrumenting OSG sites with standardized tools to gather network metrics
  - **OSG Network Service:** Gathering OSG network metrics centrally and making them available for users and applications
  - **Network Documentation:** Creating documentation for OSG user and VO managers to guide them in understanding and diagnosing network issues
- Year 2 primary components:
  - Complete deployment of perfSONAR-PS
  - Improving the modular dashboard
  - Explore extending coverage to include WLCG
  - Enable alarming and problem analysis based upon network metrics
  - Improve tools and documentation from user perspective

Replacement Prototype: MaDDash

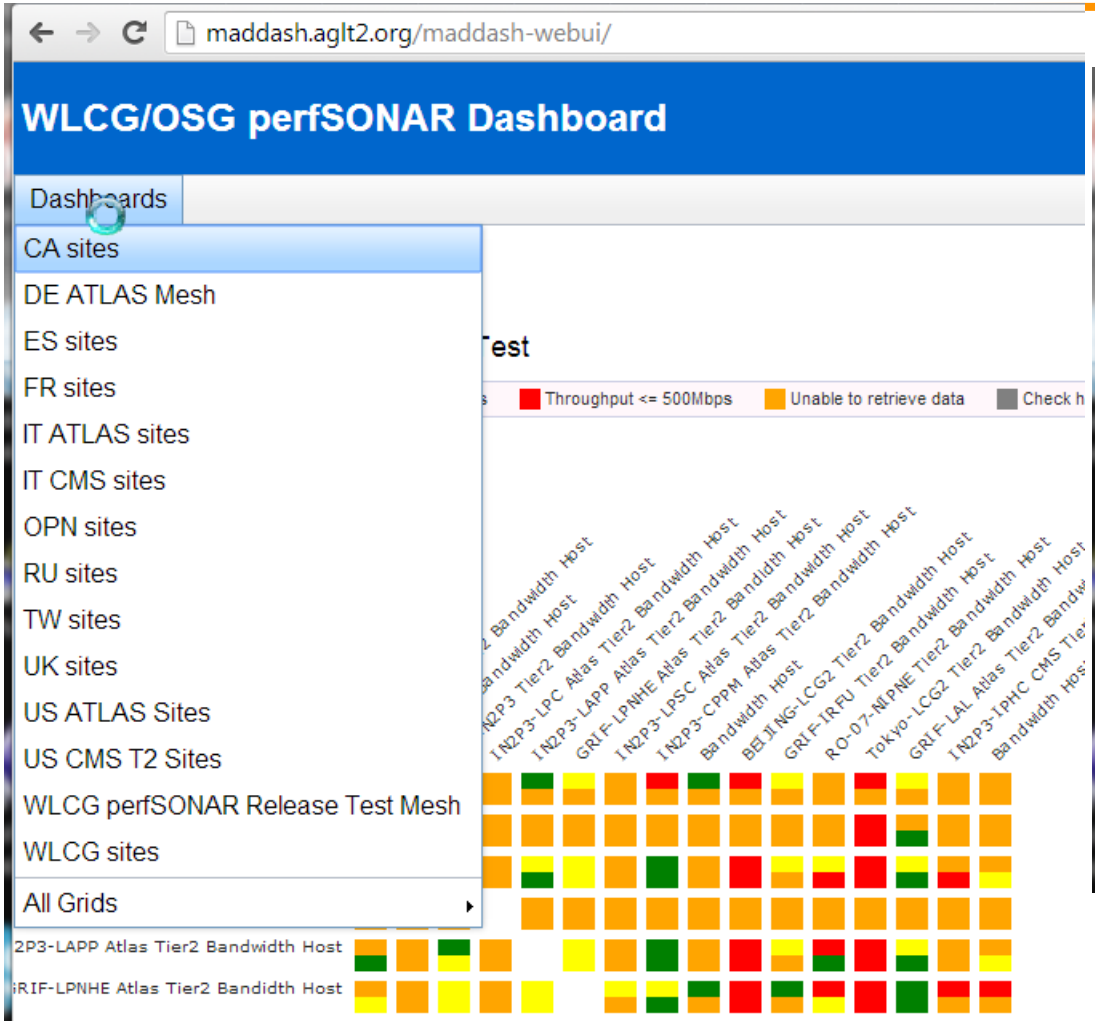

Must be migrated to OSG!

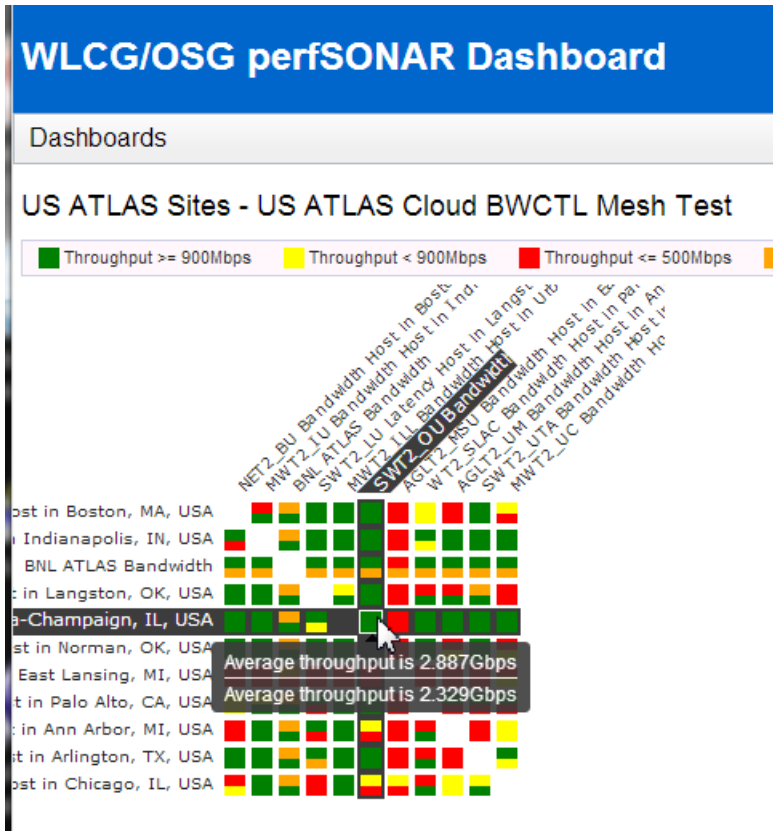

MaDDash (Monitoring and Debugging Dashboard) supported by ESnet OSG Networking

Prototype: Service Monitoring

Must be migrated to OSG!

OMD (Open Monitoring Distribution)

Integrated package over Nagios

Checks/verifies primitive services are functional

Ensures we get good network metrics

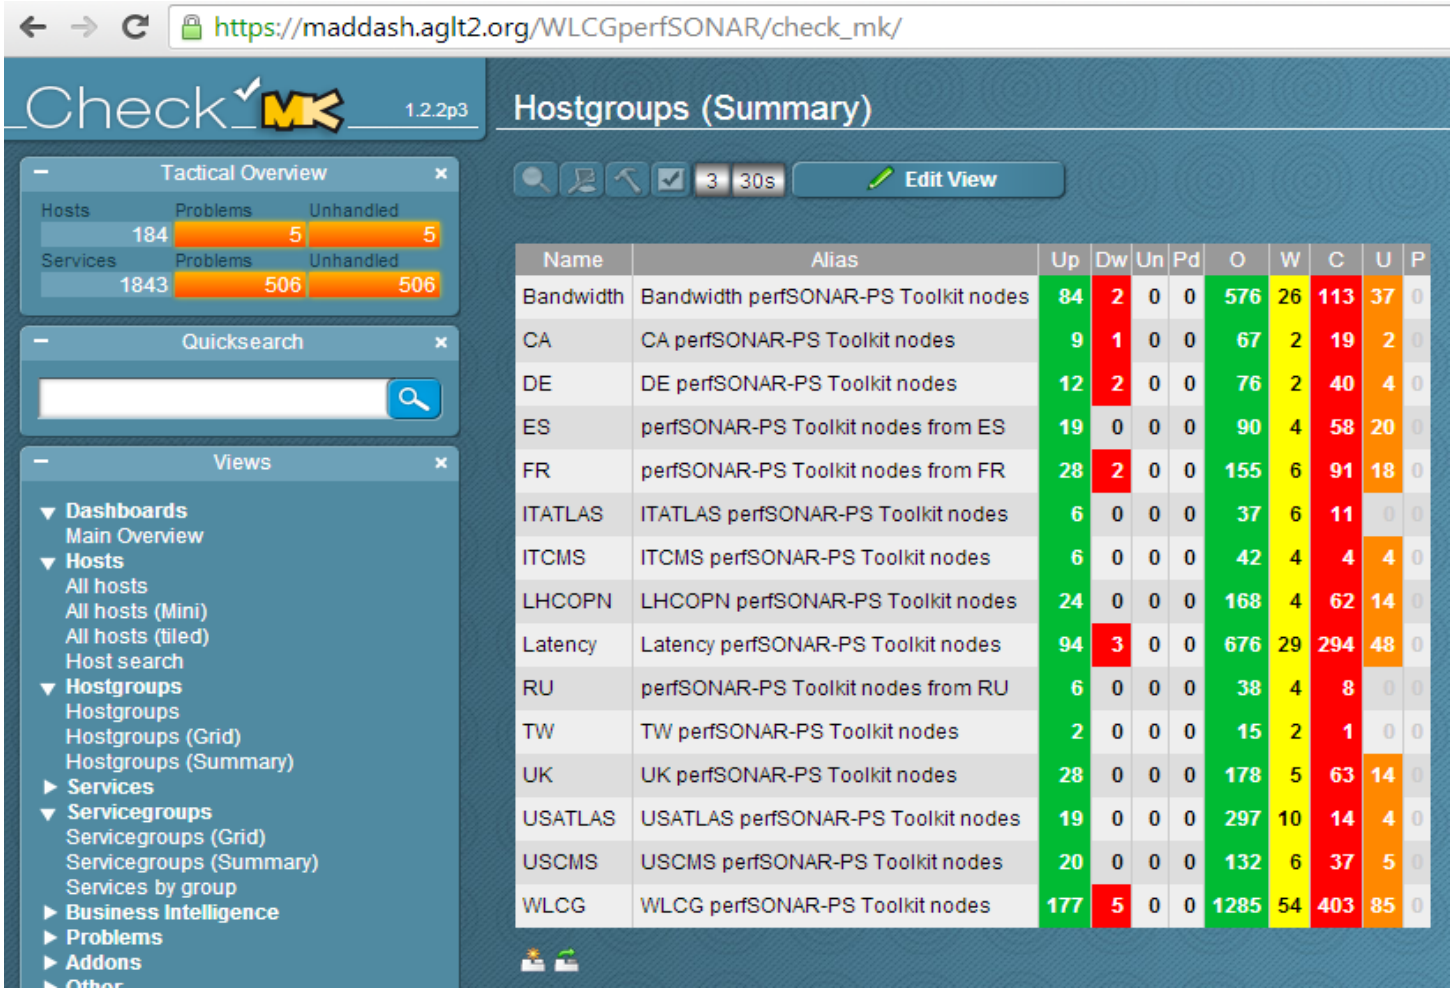

Name	Alias	Up	Dw	Un	Pd	O	W	C	U	P
Bandwidth	Bandwidth perfSONAR-PS Toolkit nodes	84	2	0	0	576	26	113	37	0
CA	CA perfSONAR-PS Toolkit nodes	9	1	0	0	67	2	19	2	0
DE	DE perfSONAR-PS Toolkit nodes	12	2	0	0	76	2	40	4	0
ES	perfSONAR-PS Toolkit nodes from ES	19	0	0	0	90	4	58	20	0
FR	perfSONAR-PS Toolkit nodes from FR	28	2	0	0	155	6	91	18	0
ITATLAS	ITATLAS perfSONAR-PS Toolkit nodes	6	0	0	0	37	6	11	0	0
ITCMS	ITCMS perfSONAR-PS Toolkit nodes	6	0	0	0	42	4	4	4	0
LHCOPN	LHCOPN perfSONAR-PS Toolkit nodes	24	0	0	0	168	4	62	14	0
Latency	Latency perfSONAR-PS Toolkit nodes	94	3	0	0	676	29	294	48	0
RU	perfSONAR-PS Toolkit nodes from RU	6	0	0	0	38	4	8	0	0
TW	TW perfSONAR-PS Toolkit nodes	2	0	0	0	15	2	1	0	0
UK	UK perfSONAR-PS Toolkit nodes	28	0	0	0	178	5	63	14	0
USATLAS	USATLAS perfSONAR-PS Toolkit nodes	19	0	0	0	297	10	14	4	0
USCMS	USCMS perfSONAR-PS Toolkit nodes	20	0	0	0	132	6	37	5	0
WLCG	WLCG perfSONAR-PS Toolkit nodes	177	5	0	0	1285	54	403	85	0

Alerting/Alarming for Network Issues

- **What most sites want is a tool that lets them know if there is a network problem (and ideally WHERE it is)**
- In year 2 we started to develop this capability for OSG sites
  - Primitive OSG perfSONAR-PS service monitoring is easy and we have Nagios-type plugins that check services
  - Much harder is deciding when network metrics gathered by perfSONAR-PS require an alert or alarm:
    - Is the change in metrics due to “normal” (heavy) network use or is there a new problem?
    - If there is a real problem, where is it located? This is critical because we should only alert someone if the problem is one they can fix
- Interesting project at Georgia Tech called Pythia (see Terena presentation https://tnc2013.terena.org/core/presentation/40)
  - Submitted new proposal NSF SI2-SSE “PuNDIT” (Pythia Network Diagnosis Infrastructure) which targets OSG/WLCG
  - Goal is to provide this needed alerting/alarming component

Network Area Near Term Goals

- OSG is strongly encouraging non-WLCG sites to deploy perfSONAR-PS toolkit instances so we can help them with network issues.
- Automating the creation of “mesh-configurations” using OIM and GOCDB registration information
- OSG production has older network datastore and monitoring in place BUT it must be merged with newer replacements.
  - Prototype services need to migrate into OSG from AGLT2
  - Must integrate new RESTful API components from perfSONAR v3.4
  - Must test API and client use-cases from OSG and WLCG
- We must evaluate the impact of monitoring and gathering network metrics for all of WLCG before committing to provide their monitoring and data aggregation.

OSG Eco-system

All OSG Usage for 12 months ending 31-March-2014

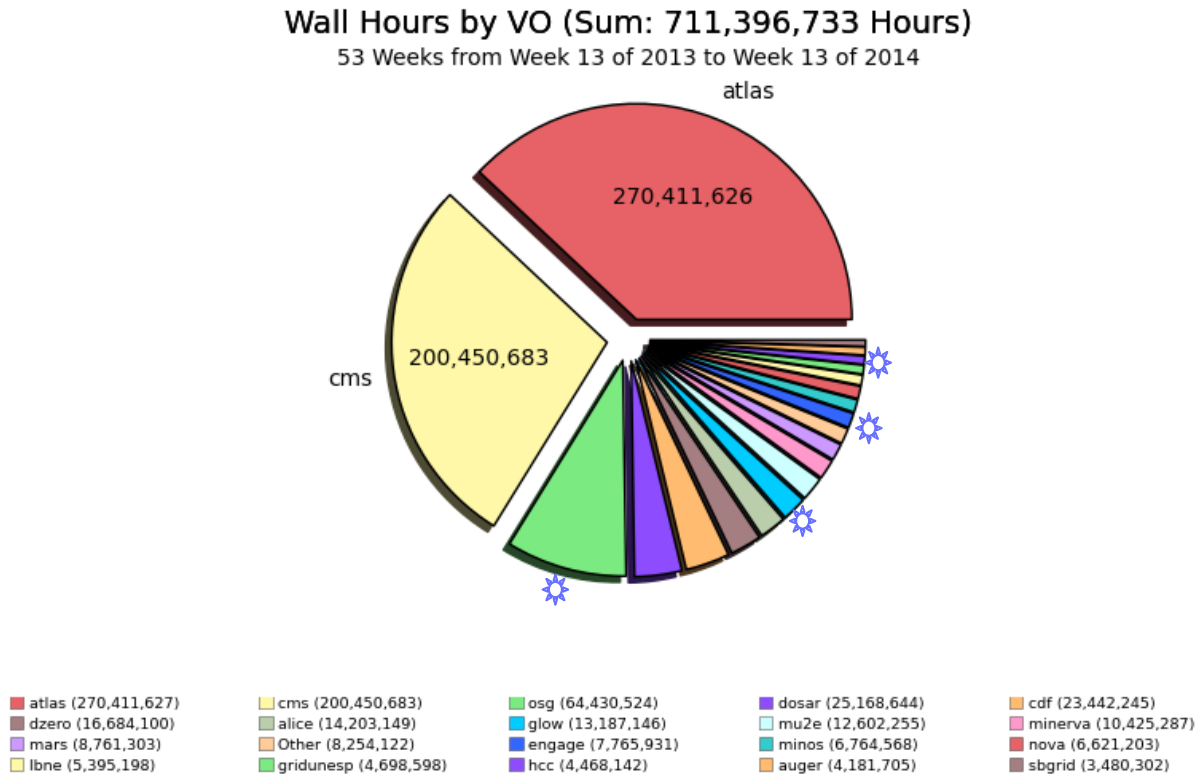

Some of these VOs access opportunistic cycles
e.g. osg, glow, engage, hcc, sbgrid

OSG Opportunistic Eco-system

Usage by “opportunistic VOs” for 12 months ending 31-March-2014

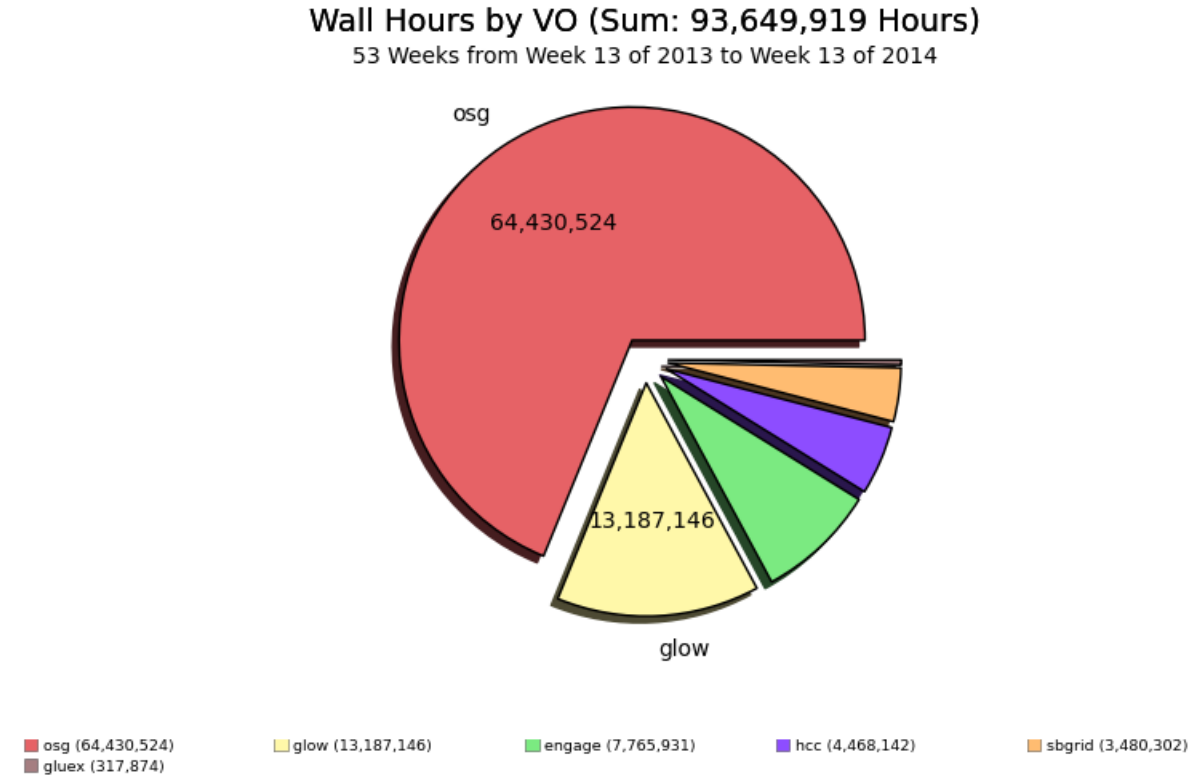

Of these, the OSG VO provides access to US researchers who are not already affiliated with an existing community in OSG

OSG VO Mission & Usage

The OSG VO does not own any computing resources and only exists to harvest unused cycles at OSG sites (Opportunistic cycles) and make them available to researchers who are not already affiliated with an OSG VO.

For the 12 months ending 31-March-2014, the OSG VO harvested 64.4M hours (from sites by using gWMS) and delivered 57.7M hours to various submit hosts to enable the computing of researchers

Submit Host	Wall Hours
OSG-XD (XSEDE and OSG Direct)**	54,694,294
UCSDgrid	1,104,882
Bakerlab	1,012,264
OSGCONNECT **	870,640
ISI	3,539
LSU	63
Total	**57,685,682**

** Core OSG Services

Access to OSG DHTC Fabric via OSG VO

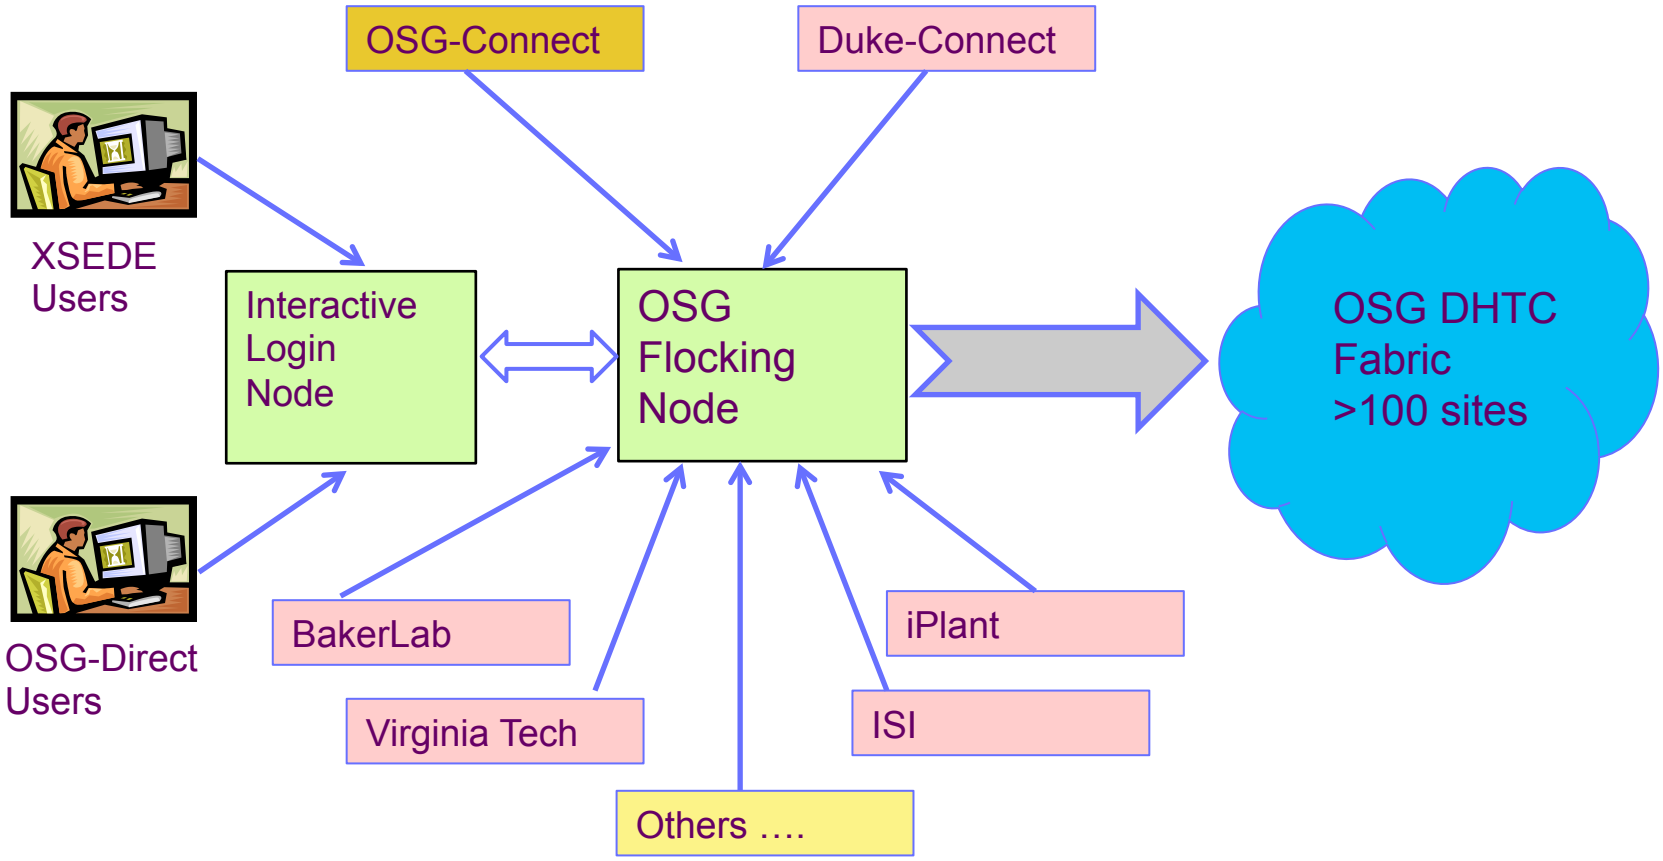

All access operates under the OSG VO using glideinWMS

Open Science Grid

OSG-Direct users April 2013 to March 2014

Project Name	PI	Institution	Field of Science	Wall Hours
Snowmass	Meenakshi Narain	Brown University	High Energy Physics	8,632,986
SPLINTER	Robert Quick	Indiana University	Medicine	4,601,962
Duke-QGP	Steffen A. Bass	Duke University	Nuclear Physics	2,543,933
ECFA	Meenakshi Narain	Brown University	High Energy Physics	1,744,646
UMich	Paul Wolberg	University of Michigan	Microbiology	1,433,598
Pheno	Stefan Hoeche	SLAC	High Energy Physics	1,108,623
RIT	P. Stanislaw Radziszowski	Rochester Institute of Technology	Computer Science	721,291
UPRRP-MR	Steven Massey	Universidad de Puerto Rico (UPRRP)	Bioinformatics	714,359
IU-GALAXY	Robert Quick	Indiana University	Bioinformatics	640,484
DetectorDesign	John Strologas	University of New Mexico	Medical Imaging	451,803
EIC	Tobias Toll	Brookhaven National Laboratory	Accelerator Physics	410,594
OSG-Staff	Chander Sehgal	Fermilab	Computer Science	43,948
DeerDisease	Lene Jung Kjaer	Southern Illinois University	Biological Sciences	28,599
SNOplus	Joshua R Klein	University of Pennsylvania	Physics - Neutrino	489
P0-LBNE	Maxim Potekhin	Brookhaven National Laboratory	Physics - Neutrino	17
BNLPET	Martin Purschke	Brookhaven National Laboratory	Medical Imaging	1
Total		**16 users**		**23,077,333**

Open Science Grid

XSEDE users April 2013 to March 2014

Project Name	PI	Institution	Field of Science	Wall Hours
TG-IBN130001	Donald Krieger	University of Pittsburgh	Biological Sciences	29,495,083
TG-PHY120014	Qaisar Shafi	University of Delaware	Physics	528,458
TG-TRA100004	Andrew Ruether	Swarthmore College	Other	444,374
TG-DMR130036	Emanuel Gull	University of Michigan	Materials Research	318,768
TG-MCB100109	Lillian Chong	University of Pittsburgh	Molecular Biosciences	264,362
TG-CHE130091	Paul Siders	University of Minnesota; Duluth	Chemistry	86,280
TG-ATM130015	Phillip Anderson	University of Texas at Dallas	Atmospheric Sciences	77,169
TG-IRI130016	Joseph Cohen	University of Massachusetts; Boston	Information; Robotics; and Intelligent Systems	70,536
TG-DMS120024	Benjamin Ong	Michigan State University	Mathematical Sciences	68,908
TG-CHE130103	Jeremy Moix	Massachusetts Institute of Technology	Chemistry	58,355
TG-ATM130009	Phillip Anderson	University of Texas at Dallas	Atmospheric Sciences	39,971
TG-MCB090163	Michael Hagan	Brandeis University	Molecular Biosciences	38,590
TG-OCE130029	Yvonne Chan	University of Hawaii; Manoa	Ocean Sciences	31,670
TG-TRA120014	Pol Llovet	Montana State University	Cross-Disciplinary Activities	19,472
TG-IBN130008	Jorden Schossau	Michigan State University	Biological Sciences	16,857
TG-MCB120070	Joseph Hargitai	Albert Einstein College of Medicine	Molecular Biosciences	378
TG-TRA120041	Hanning Chen	George Washington University	Computer and Information Science	231
TG-MCB090174	Shantenu Jha	Rutgers University	Molecular Biosciences	58
TG-PHY110015	Pran Nath	Northeastern University	Physics	37
TG-MCB130072	Robert Quick	Indiana University	Molecular Biosciences	16
TG-CCR120041	Luca Clementi	San Diego Supercomputer Center	Computer and Computation Research	12
TG-STA110014S	Nancy Wilkins-Diehr	University of California-San Diego	Other	5
Total		**22 users**		**31,559,590**

Operations Mission and Structure

The mission of OSG Operations is to maintain and support a production quality computing environment for research communities.

- Operations Support
  - Support Desk
  - Ticket Tracking
  - Community Notification and Communication
- Operations Infrastructure
  - Compute Services
  - Distributed
    - IU, FNAL, UCSD, UNL, UC

Service Levels

- Maintaining All Services at SLA Levels
  - This includes compute and support services.
  - All compute services at 99.41% Availability
    - Only missed a single monthly metric for MyOSG in July 2013
  - All critical services 99.92% Availability
    - Outage could lead to mass job failure
    - This is approximately 12 hours between June 2012 and February 2014.
  - Service Desk – No exceptions to SLA

Communication and Interoperability

- Continual Communication via Notifications, Blog Aggregation, Real Time Operational Event Tracker
- Inter-Area Communication and Coordination with Major Stakeholders
  - Bring all area coordinator together weekly for Production meeting
  - ATLAS, CMS, and Invited VOs
- Ongoing Collaboration with WLCG and EGI
  - ENMR VO fully interoperational
- Interoperability for peering infrastructures researchers
  - WLCG, XSEDE Campus Bridging, EGI-Inspire

Impact of Production Foundations

- Stable Infrastructure
- Timely Support
- Adoption of New Technologies
- Continual Communication
- Resource and Infrastructure Monitoring

Impact of Production Foundations

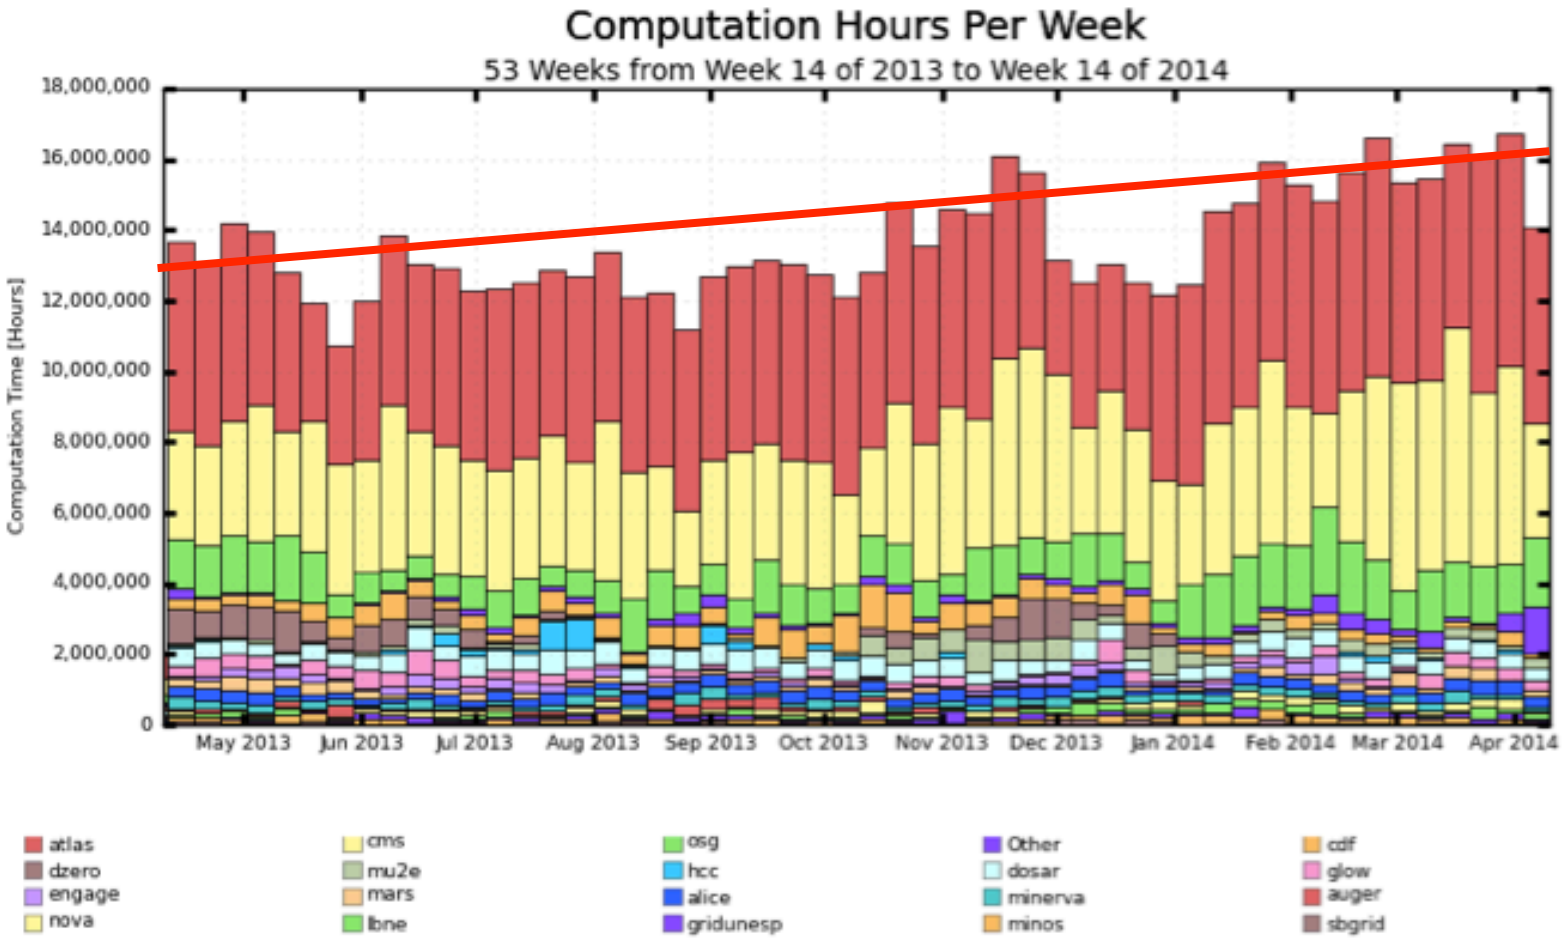

Things We Learned Yesterday

- “Researchers are people.” -Lauren
- “You have to believe in sharing.” –Miron

- So what are we doing?
  - Networks
  - Science
  - Operations
  - Technology
  - Karaoke
- But what are we really doing?

The Real Challenge of Operations

- Building a strong sense of community for users, resource suppliers, and OSG staff
  - Stable Services
  - Built in continued quasi-daily one-on-one interactions
  - Done in long term dialogues
  - You can not have a sense of community without a sense of caring.

"What should young people do with their lives today? Many things, but the most daring thing is to create stable communities…" Kurt Vonnegut

Thoughts?